

Investing in the Harness

My Thesis on Applied AI, Governance and Sovereignty

Marc Taverner, The Polaris Group · October 2026

When I started in internet infrastructure, most people assumed the value would sit with whoever owned the fastest pipe. It didn't. Bandwidth got cheaper every year, and the money moved to the people who built the things that made a connection useful and trustworthy for a particular business: payments, identity, security, the dull plumbing of who is allowed to do what. I spent the next part of my career in collaboration platforms and then in digital assets, and I watched the same pattern repeat. The raw capability commoditises. The control layer around it does not.

I believe applied AI is following that pattern, and faster than most investors are pricing in.

The model is not the asset

Foundation models have become broadly interchangeable for most business work. Per-token prices keep falling while volume keeps growing, and the gap between the best model and the third-best model on any given task is now measured in months, not years. Most serious buyers I speak to already run more than one model and choose by price and fit for each task.

That has a simple consequence for where I put money. If the model is a commodity input, a company whose value is really just the weights it trained, or the fine-tune it sits on, is renting its advantage from someone who can reprice it at will.

The durable asset is what I call the harness. It is everything built around a model to make it reliable inside one particular business. Swap the model and a good harness should keep working.

What I mean by the harness

I think about the harness as six layers.

- **Context** is the documents, history and retrieval that give an agent the knowledge of this business rather than the internet in general.
- **Policy** is what the agent may and may not do, written down in a form a machine can enforce.
- **Process** is the workflows, tools and orchestration that turn a single answer into a piece of finished work.

- **Evaluation** is the test sets, graders and traces that define what a correct answer looks like in this particular business.
- **Feedback** is how a correction made today improves the system tomorrow.
- **Identity** is who the agent is, who it acts for, and what it is entitled to touch.

These layers are not equally investable. Context, in the narrow sense of conversational memory, became a free platform feature on major platforms between September 2024 and March 2026. I would not fund it as a standalone product today, however good the team.

Evaluation and identity are a different story. The definition of a correct answer in an insurer's claims process, or a bank's onboarding flow, is specific to that organisation. It is also exactly the thing an auditor, a regulator or an insurer will ask to see when an agent gets something wrong. Identity is the question every security team is now being asked by its board: which agents are running, on whose authority, with which credentials, and who can switch them off.

What I want to fund is the governance, lineage and evaluation infrastructure that works across clouds and across model vendors, produces evidence that a regulator or auditor will actually accept, and is sold to the security and compliance buyer. One condition matters more than any other. The customer keeps ownership of its own evaluation content and test cases. That is the asset. The vendor's job is to make it usable, not to hold it hostage.

Leasehold and freehold

There is a property analogy I keep coming back to. A fine-tune on a closed frontier model is a leasehold, not a freehold. OpenAI closed self-serve fine-tuning to new customers on 7 May 2026, and from 6 January 2027 existing customers can no longer create new fine-tuning jobs either. On Azure there is no way to export fine-tuned weights, and fine-tuned deployments are deleted after fifteen days of inactivity, or twenty four hours on the developer tier. The model itself is retained and can be redeployed, but you never hold the weights. You paid for the improvement, and you still don't own it.

A company's own evaluation suite, its written policies and its curated context are different. They sit on its own systems. It holds them by possession rather than by contract, which is the only form of real ownership this market currently offers.

Whether that harness genuinely survives a change of underlying model is something I am treating as a thesis to test, not a fact to assert. I think it does, mostly. I don't know yet how much rework is involved in practice. So I am measuring it myself. I am building my own personal AI agent, a genuinely personal assistant that does my digital work while keeping my memory, context and data under my control rather than captured inside one vendor's platform. Every so often I swap the underlying model and record two things: how much rework the change needs, and whether a fixed evaluation suite catches the resulting drift before it causes a real mistake.

Who actually pays

Here is the commercial reality I have learned to respect. Individuals will not pay for privacy they can get for free, nor do the majority care about privacy beyond the utility it can deliver. For example, using a VPN to access content otherwise geolocation restricted. Roughly four in five password manager users sit on the free tier. VPN users are more willing to pay, and free use there is a fast growing minority rather than a majority, but the direction is the same. I like consumer privacy as a cause. I don't back it as a business.

Enterprises are different because they have obligations, not preferences. They have audit requirements. They have data residency law. They carry contractual liability for what an autonomous agent did on their behalf. Nobody buys a compliance control for fun; they buy it because the alternative is a fine or an uninsurable risk. That is the buyer I want the companies I back selling to.

Revealed enterprise spend on securing AI is still a small fraction of total AI spend. Some people read that as a ceiling. I read it as an early category. Cloud security, identity and endpoint protection all started as rounding errors on the budget of the thing they protected, and each became a large market once something went wrong publicly and the obligation caught up with the technology.

I also want to be honest about where the independent opportunity is narrower than the headline suggests. An enterprise that runs everything through one model vendor will simply buy controls from that vendor. The opening for an independent company is in three places. The first is enterprises running agents across several clouds and several model vendors, where no single vendor's tooling can see the whole picture. The second is enterprises with real jurisdictional exposure, where the incumbent vendor is itself part of the risk being managed. The third is any situation that needs evidence for a genuinely independent third party, such as an auditor or an insurer, who will not take a vendor's word about its own product.

That last point is the heart of it. Neutrality is the product.

What actually works for privacy today

A great deal of money is being raised on privacy technology that does not yet do what the pitch says, at least not for AI inference.

Confidential computing works. It runs a workload inside a secure hardware enclave, so that even the cloud operator cannot see the data while it is being processed, and it can prove to the customer which code ran. It is in production at the major clouds and on current GPUs, and the performance overhead has become small. Keeping sensitive context on the client side and sending only what a task needs also works, and it costs nothing.

Fully homomorphic encryption and multi-party computation are not practical for AI inference at any meaningful scale today. They are serious research fields. They are not where I would put a privacy guarantee for a production agent in 2026.

Zero-knowledge proofs are the most frequently misapplied of all. A zero-knowledge proof shows that a computation was carried out correctly. It does not let a model compute on data it cannot see. When someone pitches me zero-knowledge privacy for AI, my question is simple: who holds the unencrypted data or supplies the trust protocols while the proof is being generated? The answer usually ends the conversation.

Simple tokenisation or data masking is not enough on its own either. Published re-identification studies show it can often be reversed.

What works today is a combination. Confidential computing for workloads that need cloud scale. Client-held context, so the sensitive material never leaves the customer's control by default. Local inference on open-weight models for the genuinely sensitive fraction of the work. Everything else runs on whichever frontier model is best and cheapest that month, and a cross-vendor evaluation suite governs the whole arrangement so the customer knows when quality moves. None of that needs a breakthrough. It needs good engineering and someone neutral to stitch it together.

Europe is the long-dated opportunity

Most of what I have described is a race against the platforms' release schedules, and I am realistic about who usually wins that race. There is one opportunity that does not depend on winning it, because it is created by law and politics rather than by a product roadmap. That is European sovereignty.

AWS's European Sovereign Cloud reached general availability in January 2026. It runs as a separate cloud partition, operated through a German holding company and three German subsidiaries, with EU national managing directors and an advisory board of five EU residents. Three of those five are Amazon employees, and the customer contracting entity is still Amazon Web Services EMEA in Luxembourg, so this is corporate separation inside the Amazon group rather than independence from AWS. The first region is in Brandenburg, backed by a stated commitment of 7.8 billion euros in Germany through 2040, with named customers including EWE, the Medizinische Universität Lausitz, Sanoma Learning and, since May, SCHUFA.

That gap between the marketing and the ownership is not a reason to dismiss it. It is the argument for the neutral layer I want to fund, because a European enterprise that needs to prove where its data went cannot rely on an attestation from a subsidiary of the company it is trying to be independent of.

At the same time, Mistral AI is positioning itself as Europe's own independent sovereign AI champion while being a launch partner on that same American hyperscaler's sovereign cloud,

and while holding a multibillion dollar partnership with Microsoft announced in July 2026 under which Microsoft draws on Mistral's European GPU capacity for its own cloud. Its largest outside shareholder is ASML. I don't see a contradiction there. I see a market still forming its coalitions, which is exactly when the infrastructure that sits between the parties gets built and gets valuable.

The legal pressure is durable. The EU Data Act's switching provisions have applied since September 2025. What lands on 12 January 2027 is narrower and sharper: switching charges go to zero altogether, and because the regulation defines switching charges to include data egress, egress for a mandated move goes with them. That makes it materially easier for a European enterprise to move workloads and materially harder for a vendor to lock them in. Egress for running two providers in parallel survives, at cost.

The EU AI Act creates real obligations around fine-tuning and deployment that a vendor based outside Europe cannot simply contract its way around. Fine-tuning can make you the provider of the model once the compute you spend on it passes roughly a third of the original model's training compute. Deployment can do the same if you rebrand a system, substantially modify it, or change what it is for. Those roles are allocated by law, not by contract.

These obligations do not get commoditised by a feature launch. They do move, though, and in 2026 they moved the wrong way for anyone underwriting enforcement. In July 2026 the Union deferred the high risk obligations by sixteen months under its simplification agenda, to December 2027 and August 2028. It left the general purpose AI obligations untouched. I am underwriting the general purpose limb, which has applied since August 2025, not the high risk timetable.

I don't want to fund another sovereign cloud. What I want to fund is the governance, lineage and evaluation layer that lets a European enterprise prove to a regulator that its agents ran where they were supposed to, under the law that actually applies, across infrastructure that mixes American hyperscalers, European model providers and open-weight models running on its own hardware. Proof is the product.

Two shapes of company

I seek to back and work with two different shapes of company.

The first shape is near-term. These are companies built to be acquired by identity and security incumbents who need the capability quickly and cannot build it in time. Okta, CrowdStrike, Cisco, SailPoint and Cyera are the kind of buyer I have in mind.

Recent acquisitions of focused security and identity companies have mostly landed between one hundred million and one billion dollars. The median priced deal in these categories sits nearer three hundred million than a billion, and roughly nine in ten acquisitions close on undisclosed terms, which usually means smaller rather than larger.

The second shape is longer-duration. These are companies whose advantage is genuine neutrality across clouds and model vendors, or genuine jurisdictional specificity, most often European. They do not need to win a sprint against a hyperscaler, because the thing protecting them is law and neutrality rather than a feature lead. They take longer and need more patient capital. They are also the ones that can become durable independent businesses.

The companies I find interesting

The first area is agent governance, identity and audit, and it has moved fast. Onyx Security is building controls for how AI agents operate inside the enterprise, so a security team can see and govern what those agents actually do. Keycard is working on identity and access built specifically for agents, the question of who an agent is and what it may touch, designed in from the start rather than retrofitted from human login systems. AIUC approaches the same problem from the insurer's side, setting a standard for AI agent risk and tying it to underwriting, which is precisely the independent third party I described earlier. GitGuardian built its name on detecting leaked secrets in code and has become a natural home for the credentials sprawl that agents create.

Two things that were on my list here have already gone, and both are the pattern rather than exceptions to it. Oasis Security came out of the non-human identity problem, the service accounts, keys and tokens that now outnumber human users in most organisations, and extended that into managing the access AI agents are granted. Cyera bought it in September 2026 for about a billion dollars. Cisco bought Galileo, an evaluation and observability company I would otherwise have named below, in May 2026 and folded it into Splunk.

These companies sit on an obligation rather than a preference, each is most useful when it sees across vendors, and each produces evidence someone other than the vendor will rely on. They have been funded at seed through Series C. None of them yet discloses meaningful revenue publicly. I say that as an observation, not a criticism. It tells me the category is still being underwritten on thesis and on team, which is when the best entry points exist and when judgement matters most.

Cyera and Cisco are consolidating this exact space quickly, Cyera having bought Otterize and then Oasis, Cisco having bought Astrix, WideField and Galileo. That is good for exit prices and bad for anyone who wants to build the second, longer-duration shape here..

Evaluation and observability infrastructure is the second area. Braintrust is a good example of what I mean: tooling that lets a company define what a correct answer looks like in its own business, trace every agent decision, and keep that evaluation suite as its own asset regardless of which model sits underneath. This is the layer that makes the leasehold and freehold distinction real for a customer.

Digital identity wallet middleware is the third. Every member state is required to offer an EU Digital Identity Wallet by 24 December 2026. As I write, one has. Denmark went into production in June. Germany, France and Italy are in public sandbox, no wallet is yet fully certified under the EU scheme, and Sweden is not expecting one before 2028. Someone has to connect those wallets to the agents acting on a person's or a company's behalf, and the gap between a binding deadline and a single live implementation is the opportunity rather than an argument against it. Procivis, the Swiss firm building wallet and credential infrastructure already used in government rollouts, shows the shape of the thing I mean, though it has been majority owned by Orell Füssli since 2021, so it is an illustration rather than a target. What I am looking for is the plumbing that lets an agent present verified credentials and act within a mandate the human has actually granted.

Post-quantum crypto-agility is the fourth. Every signed credential, every agent identity and every audit log I have described rests on cryptography that will need replacing over the next decade at least & maybe sooner. PQShield, the Oxford firm selling post-quantum tooling to enterprises and silicon vendors, is an interesting example: not a new algorithm, but the ability for an enterprise to find, inventory and swap its cryptography without rebuilding its systems. Regulators have already set migration timelines, so this is an obligation with a date attached.

Confidential computing and attested inference is the fifth. Anjuna Security is building the kind of product I described in the privacy section: running workloads, AI inference included, inside hardware-attested enclaves across any cloud, and giving the customer a cryptographic receipt showing which model ran, on which hardware, against which data. That receipt is exactly the sort of evidence an auditor can accept.

European sovereign AI infrastructure is the sixth. OVHcloud is the right shape not because it competes with the hyperscalers on scale, but because it gives European enterprises a credible, jurisdictionally clean option for the part of their work that has to stay under European law. Its agreement to acquire Dragon LLM, signed in March 2026, and the sales organisation it has built around sovereign demand shows that it is building the plumbing, not just the marketing.

Across all six, the common thread is the same. The company makes the harness portable, provable or protected, and it sells to someone who is obliged to buy.

What I'm not too interested in

The following examples I find less appealing to engage with;

- Model companies, and anything whose value is really just the underlying weights.
- Anything relying on fully homomorphic encryption, multi-party computation or zero-knowledge proofs as the actual privacy mechanism for AI inference. I've no evidence they work at this scale yet.
- Anything whose entire differentiation is a feature already on a hyperscaler's public roadmap.

- Any company that wants to own the customer's evaluation data and test cases rather than letting the customer keep them. That is lock-in dressed up as a sovereignty pitch, and I won't back it.

What would change my mind

There are two open questions I can't yet answer, and they are the two that would move me most.

The first is portability in practice. I have not yet found a single named enterprise that has publicly credited the portability of its own evaluation harness with letting it switch its primary AI vendor. I believe it is happening quietly. I would like to see it on the record. If, over the next year or so, large buyers start citing their own evaluation suites as the reason a model switch took weeks rather than quarters, the freehold argument becomes a proven commercial one. If switches keep happening and nobody credits the harness, I will need to rethink how much of this value really sits with the customer.

The second is who captures the spend. I don't yet have good data on whether enterprises buy this kind of governance tooling from independent companies, or simply absorb it into contracts they already hold with their incumbent security vendors. If independent companies keep winning standalone budget lines, the second, longer-duration shape gets much more attractive. Does the spend fold into existing Okta, CrowdStrike or Microsoft agreements?

A third has been added by events rather than by argument. The European regulatory timetable moved in 2026, and it moved toward simplification rather than enforcement. If the high risk deferral turns out to be the first of several, and the general purpose obligations are softened in the same way, the European limb of this thesis weakens considerably. I am watching the digital omnibus files, which were still unresolved in Council as I write.

Where I have landed

My underlying convictions haven't changed since I first started writing about this. I care about portability. I care about sovereignty, at the personal level and at the national level. I care about privacy, and I want agentic AI that is genuinely useful and genuinely trustworthy, not one or the other.

What has changed is how precisely I can say where those convictions meet a cheque I'm willing to write. I back companies that own the harness rather than the weights, that sell to an obligation rather than a preference, and that compete on neutrality rather than secrecy. I back them in Europe where the law gives them a long runway, and anywhere an independent party needs evidence a vendor can't provide about itself.

If you're building one of those companies, or investing in them, or writing the rules they will have to prove they follow, I'd like to talk.

Evidence notes

This edition was checked against primary sources on 22 and 23 October 2026. Data and content is accurate to then as best I have been able to check. I am human, I may make mistakes & I may have thoughts and opinions you don't agree with. These are likely driven by my own personal experiences and humanity.

- **OpenAI fine-tuning.** OpenAI deprecations page, "Update to OpenAI's self-serve fine-tuning", announced 7 May 2026. Three stages: 7 May 2026, 2 July 2026, 6 January 2027. developers.openai.com/api/docs/deprecations
- **Azure fine-tuned deployments.** Microsoft Learn, Azure OpenAI fine-tuning deployment and cost management pages. Fifteen day inactivity deletion; twenty four hours on the Developer tier. No Microsoft document states that weights cannot be exported; there is simply no mechanism.
- **Password manager and VPN free tiers.** Security.org password manager annual report, 79 per cent on free apps, fieldwork October 2024. Security.org VPN consumer report, 43 per cent on free services, fieldwork August 2026.
- **AWS European Sovereign Cloud.** AWS News Blog 14 January 2026 and AWS press release, Potsdam, 15 January 2026. Governance and entity structure from aws.eu/governance and the AWS overview whitepaper. Investment figure of 7.8bn euros through 2040 from the AWS announcement of 15 May 2024, repeated January 2026.
- **Mistral.** AWS launch partner list, 15 January 2026. Microsoft and Mistral expanded partnership, 21 July 2026. ASML Series C, 9 September 2025, approximately 11 per cent fully diluted.
- **EU Data Act.** Regulation (EU) 2023/2854. In force 11 January 2024, applicable from 12 September 2025. Article 29(1) ends switching charges from 12 January 2027; Article 2(36) defines switching charges to include data egress. Recital 99 preserves egress charges for in parallel use, at cost.
- **EU AI Act.** Regulation (EU) 2024/1689, Articles 2, 25 and 113. Commission guidelines on the scope of GPAI obligations, C(2025) 5045 final, 18 July 2025, for the one third compute threshold. Regulation (EU) 2026/1744, in force 27 July 2026, for the deferral of Chapter III to 2 December 2027 and 2 August 2028.
- **EU Digital Identity Wallet.** Article 5a of Regulation (EU) No 910/2014 as amended by Regulation (EU) 2024/1183. Deadline of 24 December 2026, being 24 months from the entry into force of the first implementing regulations on 24 December 2024. Rollout status as at October 2026.
- **Acquisitions.** Cyera and Oasis Security, announced 28 July 2026, completed 3 September 2026, approximately \$1bn, reported. Cisco and Galileo, announced 9 April 2026, closed 22 May 2026, terms undisclosed. Google and Wiz, \$32bn, closed 11 March 2026. Palo Alto Networks and CyberArk, \$25bn, closed February 2026. Median and

disclosure-rate figures from SecurityWeek annual M&A reviews for 2024 and 2025 and from H1 2026 market reports.

- **GitGuardian.** Series C, \$50m, February 2026, led by Insight Partners.
- **Procivis.** Majority owned by Orell Füssli since September 2021, following an initial stake in spring 2020.
- **OVHcloud and Dragon LLM.** Binding agreement announced 25 March 2026. No completion announcement had been published as at October 2026.
- **Platform memory.** Free tier availability: ChatGPT September 2024, Meta AI January 2025, Gemini March 2025, Copilot April 2025, Grok around April 2025, Claude around March 2026. The claim holds for conversational memory; richer retrieval features remain subscriber-gated on several platforms.